

Removal of Drift From Frequency Stability Measurements

C. A. Greenhall

Communications Systems Research Section

This article gives a method of estimating frequency drift rate and removing its effect from Allan variance plots. When tried on a test of hydrogen masers, the method gives consistent results. An error in the previous Allan variance computation algorithm is corrected.

I. Drift Removal—Before and After

Imagine a frequency standard whose only problem is a steady frequency drift. Its phase error in radians has the form

$$\phi(t) = 2\pi f_0 \left(a + bt + \frac{1}{2} ct^2 \right) \quad (1)$$

where a , b , and c are constants, and f_0 is the nominal frequency of the oscillation. The drift rate of the relative frequency error $\Delta f/f_0 = \dot{\phi}(t)/(2\pi f_0)$ is c per second. The two-sample Allan deviation, the usual measure of $\Delta f/f_0$, is

$$\sigma_y(\tau) = \frac{|c|\tau}{\sqrt{2}} \quad (2)$$

for this simple case. (In this report, Allan *deviation* is the square root of Allan *variance*.)

Now observe the behavior of three hydrogen masers, called DSN1, DSN2, and DSN3, which were tested at the JPL Interim Frequency Standard Test Facility for eight days at the end of 1980. Figure 1, a rough rendering of the frequency

strip charts, shows $\Delta f/f_0$ vs calendar date for the three possible pairs of masers beating against each other. *The three curves have different $\Delta f/f_0$ scales.* In particular, the scale of the DSN2-DSN3 curve is expanded relative to the others because it does not exhibit the drift that dominates the other curves. The thickness of the DSN2-DSN3 curve is just a way of showing the size of the rapid (3-minute average) fluctuations of $\Delta f/f_0$. It is a good bet that DSN1 was drifting by itself at a rate about -6×10^{-19} per second, or -5×10^{-14} per day.

Figure 2 shows what this drift does to the Allan deviation. The usual $\sigma_y(\tau)$ is given by the "gross" curves, which, for the two pairs containing DSN1, become straight lines with slope one for the larger τ . The dashed lines show the estimated drift component, Eq. (2), where the estimate of the drift rate c is computed by a method explained below. The actual estimates of c are

$$6.15 \times 10^{-19} / \text{s} \quad \text{DSN2-DSN1}$$

$$-6.32 \times 10^{-19} / \text{s} \quad \text{DSN1-DSN3}$$

$$-4.13 \times 10^{-21} / \text{s} \quad \text{DSN2-DSN3}$$

all of which have standard deviation 3.6×10^{-20} . Evidently, DSN2-DSN3 has negligible drift.

The “net” curves in Fig. 2 show what happens when the estimated drift function $(2\pi f_0) (ct^2/2)$ is subtracted (in effect) from the phase data. The net Allan deviations, $\sigma_{y0}(\tau)$, for the DSN1 pairs look like the gross Allan deviation of DSN2-DSN3. All three curves have slope 0.77 for $\tau > 10^4$ s. The effects of the random phase fluctuations, formerly masked by the drift, can now be seen.

II. Method of Drift Estimation and Removal

A. Quantities to be Estimated

It is convenient to work with the function

$$x(t) = \phi(t)/(2\pi f_0)$$

where $\phi(t)$ is the phase difference of the pair of oscillators being tested. The underlying assumption is that $x(t)$ is a mean-continuous stochastic process whose second differences

$$\Delta_\tau^2 x(t) = x(t) - 2x(t - \tau) + x(t - 2\tau)$$

are stationary for each τ . A deterministic example is given by Eq. (1); its second differences

$$\Delta_\tau^2 \left(\frac{1}{2} ct^2 \right) = c\tau^2$$

are constant. In fact, it is true in general that any such process can be written

$$x(t) = \frac{1}{2} ct^2 + x_0(t) \quad (3)$$

where c is a constant, and the second differences of $x_0(t)$ have mean zero (Ref. 1). This decomposes the phase into a pure frequency drift term plus random fluctuations. (The term $x_0(t)$ might contain an $a + bt$ component, which goes away when second differences are taken. Anyhow, we do not care about constant phase and frequency offsets.) Our goal is to perform this decomposition on experimental phase data.

The usual Allan variance, called *gross AV* here, is defined by

$$\sigma_y^2(\tau) = \frac{1}{2\tau^2} E [\Delta_\tau^2 x(t)]^2 \quad (4)$$

and the *net AV* is defined by

$$\sigma_{y0}^2(\tau) = \frac{1}{2\tau^2} E [\Delta_\tau^2 x_0(t)]^2 \quad (5)$$

Since the second differences of $x_0(t)$ have mean 0, Eq. (3) gives

$$c = \frac{1}{\tau^2} E \Delta_\tau^2 x(t) \quad (6)$$

$$\sigma_{y0}^2(\tau) = \sigma_y^2(\tau) - \frac{1}{2} c^2 \tau^2 \quad (7)$$

We want to estimate $\sigma_y(\tau)$, c , and $\sigma_{y0}(\tau)$.

B. The Estimators

Let $x(t)$ be given for $0 \leq t \leq T$. We shall need the first four moments of the second differences. Define

$$m_i(\tau) = \frac{1}{r} \sum_{j=2}^{r+1} [\Delta_\tau^2 x(j\tau)]^i$$

The integer r , which depends on τ , is the available number of second-difference samples. One can define it by saying that $(r+1)\tau$ is the largest multiple of τ that does not exceed T .

The usual estimator of gross *AV* is the time average

$$S^2(\tau) = \frac{1}{2\tau^2} m_2(\tau) \quad (8)$$

How shall we estimate c ? Equation (6) suggests the unbiased estimator

$$\hat{c}(\tau) = \frac{1}{\tau^2} m_1(\tau) \quad (9)$$

Because a second difference is the difference of first differences, the implied summation in Eq. (9) telescopes, leaving us with

$$\hat{c}(\tau, \tau') = \frac{1}{\tau'} \left[\frac{\Delta_\tau x(\tau + \tau')}{\tau} - \frac{\Delta_\tau x(\tau)}{\tau} \right] \quad (10)$$

where $\tau' = r\tau$. The notation is expanded because Eq. (10) is more general than Eq. (9), in that τ' does not need to be an

integer multiple of τ . The interpretation is that the average drift rate equals average frequency near the end of the record, minus average frequency near the beginning, divided by the length of the record (actually, by τ').

We wish to select just one estimator of c for the given record length T . To do this, we might minimize the variance of Eq. (9) or Eq. (10) over τ , where $\tau' = T - \tau$ in Eq. (10). This cannot be done in advance without knowing the spectrum $S_x(f)$ of $x(t)$. Since $S_x(f)$ determines Allan variance, we are asking for the outcome of our measurements before we do them. To escape this trap, we appeal to the past — a measurement made by Sward (Ref. 2) on hydrogen masers. His work gives the one-sided spectral density

$$S_x(f) = \frac{h_{-1}}{(2\pi)^2 f^3} \quad (\text{flicker frequency modulation})$$

$$+ \frac{h_1}{(2\pi)^2 f} \quad (\text{flicker phase modulation})$$

where

$$h_{-1} = 3.5 \times 10^{-29}, h_1 = 1.6 \times 10^{-25} \text{ s}^2$$

and the second term is cut off at $f = 10^5$ Hz.

It turns out that the variance of $\hat{c}(\tau, \tau')$ can be read from formulas in Ref. 3, pp. 42-47, for different types of phase noise, including flicker FM and PM. Using the Sward spectrum, we find, for $\tau < 15$ s, that the flicker PM part of $\text{Var } \hat{c}(\tau, \tau')$ is dominant. For $\tau > 15$ s, the flicker FM part takes over. Furthermore, as τ increases beyond 15 s, the variance becomes smaller than for any $\tau < 15$ s. Hence, we need only consider the flicker FM contribution to $\text{Var } \hat{c}(\tau, \tau')$, which is

$$\frac{h_{-1}}{\tau'^2} [(r+1)^2 \ln(r+1) - 2r^2 \ln r + (r-1)^2 \ln(r-1)] \quad (11)$$

Let us state the result of minimizing this.

Assume that the normalized phase error $x(t)$ consists of flicker FM plus a constant frequency drift term ($ct^2/2$). Then Eq. (10) gives a family of unbiased estimators of c . Let $T = \tau + \tau'$ be fixed. Then

$$\min_{\tau} \text{Var } \hat{c}(\tau, T - \tau) = \frac{8.94 h_{-1}}{T^2} \quad (12)$$

The minimum is achieved for $r = (T - \tau)/\tau = 5.29$.

In other words, we should use a τ that is about one-sixth of the record length T . Notice that the variance is like $1/T^2$ instead of $1/T$. This happens because the second differences of $x(t)$ have less power than white noise near zero frequency.

The minimum is broad enough to allow considerable departures from it. Although Eq. (10) is simple, the details of the data processing make it expedient to revert to the summed form Eq. (9). Moments of the second differences are accumulated only for a certain small set of τ , and the actual estimator of c is

$$C = \hat{c}(\tau_c)$$

where τ_c is the largest available τ such that $r_c = r(\tau_c)$ is at least 6. Of course, $\tau_c + \tau'$ is usually less than T . For the τ -set actually used, r_c falls between 6 and 16.

We can estimate $\text{Var } C$ from Eq. (12) (or Eq. (11) if τ_c is not quite optimal) if we have a value for the flicker FM constant h_{-1} , which satisfies

$$h_{-1} = \frac{1}{\ln 4} \sigma_{y0}^2 \quad (13)$$

(Ref. 3). This leads to the next goal, the estimation of $\sigma_{y0}^2(\tau)$. In view of Eq. (7), one might use the estimate

$$S^2(\tau) - \frac{1}{2} C^2 \tau^2$$

The problem with this is that it can be negative. We prefer to start from Eq. (5). Given the data $x(t)$ and the estimate C of c , an estimate of the "net data" $x_0(t)$ is $x(t) - (Ct^2/2)$ (except for a polynomial $a + bt$). Then, an estimate of $\Delta_r^2 x_0(t)$ is $\Delta_r^2 x(t) - Ct^2$. This leads to our estimator of choice,

$$S_0^2(\tau) = \frac{1}{2\tau^2 r} \sum_{j=2}^{r+1} [\Delta_r^2 x(j\tau) - Ct^2]^2 \quad (14)$$

$$= S^2(\tau) - C m_1(\tau) + \frac{1}{2} C^2 \tau^2$$

for the net $AV \sigma_{y0}^2(\tau)$.

Since $C\tau_c^2 = m_1(\tau_c)$, $S_0^2(\tau_c)$ is just $(1/(2\tau_c^2))$ times the variance of the sequence of second τ_c -differences of $x(t)$. One can now estimate h_{-1} from Eq. (13) by using $S_0^2(\tau_c)$ for σ_{y0}^2 .

For the largest useful τ , corresponding to $r = 5$ or less, an anomaly may appear. Either the estimated net AV or the estimated drift contribution $C^2\tau^2/2$ can come out greater than the estimated gross AV . This should not be alarming. If net AV is greater than gross AV , the gross AV should still fall within the error bar of the net AV . This error bar is the next topic.

The net AV estimate $S_0^2(\tau)$ is the average of the numbers

$$u_j = \frac{1}{2\tau^2} [\Delta_\tau^2 x(j\tau) - C\tau^2]^2$$

A computation with Gaussian flicker FM shows that the sampled process

$$[\Delta_\tau^2 x_0(j\tau)]^2$$

is almost white from zero frequency to the Nyquist frequency, even though the spectral density of $\Delta_\tau^2 x_0$ vanishes at zero frequency. Thus, it is reasonable to use the sample variance of the u_j for estimating the error in the mean. Our one-sigma error estimate for $S_0^2(\tau)$ is δ , given by

$$\begin{aligned} \delta^2 &= \frac{1}{(r-1)r} \sum_{j=2}^{r+1} [u_j - S_0^2(\tau)]^2 \\ &= \frac{1}{4\tau^4 (r-1)} [m_4 - 4Dm_3 + 4D^2m_2 - (2Dm_1 - m_2)^2] \end{aligned} \quad (15)$$

where $D = C\tau^2$, and the m_i are the moments $m_i(\tau)$ defined above. Finally, an error bar for $S_0(\tau)$ is $[S_0^2(\tau) - \delta]^{1/2}$ to $[S_0^2(\tau) + \delta]^{1/2}$. If $\delta > S_0^2(\tau)$, then the first number is replaced by zero.

III. Correction of an Error

In the JPL frequency stability test setup, the phase data are written on magnetic tape. Later, they are processed into Allan variance by an offline computer program. The previous version of this program contains an error. To explain it, fix a τ , and let

$$\xi(t) = \frac{1}{\tau\sqrt{2}} \frac{\Delta_\tau^2 \phi(t)}{2\pi f_0} = \frac{\Delta_\tau^2 x(t)}{\tau\sqrt{2}}$$

be a scaled version of the second difference of phase. According to Eq. (8), the usual estimator of $\sigma_y(\tau)$ is

$$S(\tau) = \left[\frac{1}{r} \sum_{j=2}^{r+1} \xi^2(j\tau) \right]^{1/2} \quad (16)$$

a discrete rms time average of ξ , just as

$$\sigma_y(\tau) = [E \xi^2(t)]^{1/2} = \|\xi(t)\|_2$$

is the rms ensemble average, or L^2 norm, of the random variable $\xi(t)$.

The previous Allan variance program used

$$S_1(\tau) = \frac{1}{r} \sum_{j=2}^{r+1} |\xi(j\tau)| \quad (17)$$

to estimate the Allan deviation. Equation (17) is an unbiased estimator, not of $\sigma_y(\tau)$, but of

$$E |\xi(t)| = \|\xi(t)\|_1$$

the L^1 norm of $\xi(t)$. One can judge the size of this error by assuming that $\xi(t)$ is Gaussian with mean zero (no drift!), in which case

$$\frac{\|\xi(t)\|_1}{\|\xi(t)\|_2} = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} |x| e^{-\frac{1}{2}x^2} dx = (2/\pi)^{1/2} = 0.7979$$

We compared the old values $S_1(\tau)$ and the new values $S(\tau)$ from another stability test run, which again measured three oscillator pairs. Figure 3 shows the comparison for one pair, a cesium standard and a hydrogen maser. For all three pairs, we computed the average of the ratios $S_1(\tau)/S(\tau)$. A ratio was included only if the standard deviation of $S_1(\tau)$ was less than 3 percent. The three averages, with error estimates, are

$$0.7992 \pm 0.0019$$

$$0.7905 \pm 0.0008$$

$$0.7968 \pm 0.0015$$

At least for this purpose, the Gaussian hypothesis seems justified. When drift is negligible, the old values for $\sigma_y(\tau)$ are 20 percent too low. This explains part of the difference between the JPL system and the Hewlett-Packard HP5390A frequency stability measurement system (Ref. 4), which does its own Allan variance computation.

It was necessary to correct this error before a drift-removal algorithm could successfully be installed in the JPL Allan variance program. Use of the L^2 norm makes it possible to remove the drift via Eq. (14) in one pass through the phase data.

IV. Concluding Remarks

The drift removal method given above yields consistent results on the data from one frequency stability test run, in that the Allan deviation curves of the three hydrogen maser pairs look almost the same after the drift is removed by the analysis program. Although the drift estimation method assumes flicker FM, the actual Allan deviation plots do not become level at the higher τ . We need more experience with the method before we can judge its robustness with respect to the flicker FM assumption. Perhaps one can find a method that tailors itself to the actual oscillator behavior. In the meantime, the present method appears to give useful results.

References

1. Yaglom, A. M., "Correlation Theory of Processes With Random Stationary nth Increments," *American Math. Society Translations*, Ser. 2, Vol. 8, pp. 87-141, 1958.
2. Sward, A., "Measurement of the Power Spectral Density of Phase of the Hydrogen Maser," *JPL Quarterly Technical Review*, Vol. 1, No. 4, pp. 30-33, 1972.
3. Barnes, J. A., et al., *Characterization of Frequency Stability*, NBS Technical Note 394, National Bureau of Standards, Washington, D.C., 1970.
4. Sosa, E. N., "Station Stability Measurement," *TDA Progress Report 42-62*, pp. 35-42, Jet Propulsion Laboratory, Pasadena, Calif., 1981.

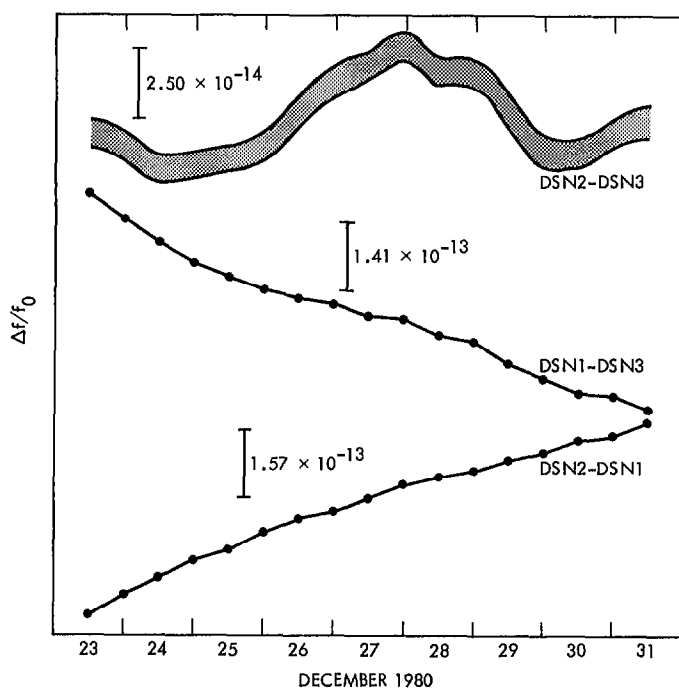


Fig. 1. A plot of pairwise $\Delta f/f_0$ vs. time for a set of three hydrogen masers, called DSN1, DSN2, and DSN3. The $\Delta f/f_0$ scales are all different

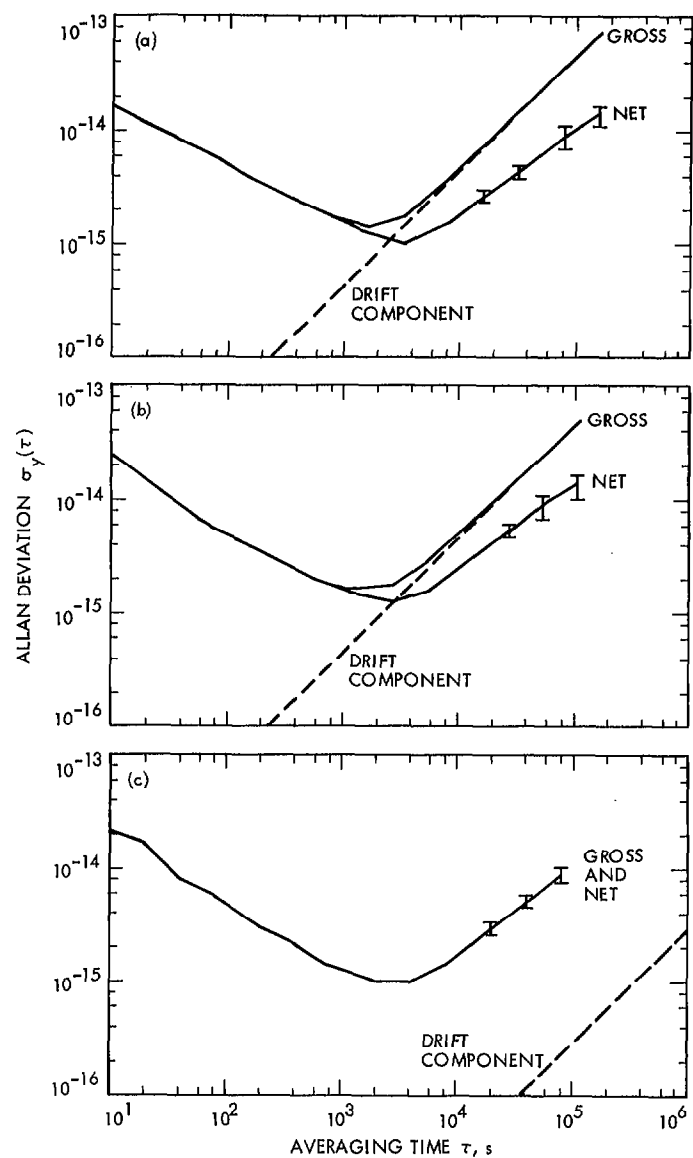


Fig. 2. Pairwise Allan deviation of three hydrogen masers, before and after removal of drift from the measurements: (a) DSN2-DSN1; (b) DSN1-DSN3; (c) DSN2-DSN3

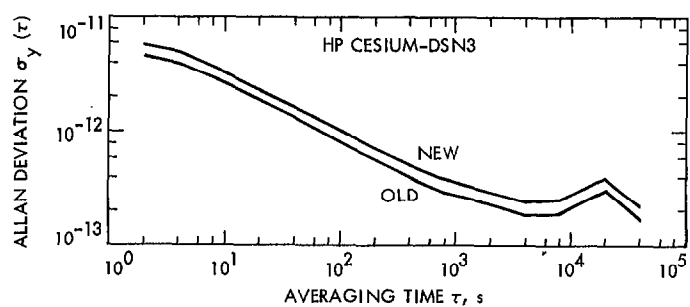


Fig. 3 Allan deviation of a cesium-hydrogen maser pair, as computed by the old (incorrect) algorithm and the new (correct) algorithm. The old results are 20 percent too low.